

Functional programming framework for *GRworkbench*

Andrew J. Moylan, Susan M. Scott and Antony C. Searle*
Centre for Gravitational Physics,
Department of Physics, Faculty of Science,
The Australian National University,
Canberra ACT 0200, Australia.

February 2, 2008

Abstract

The software tool *GRworkbench* is an ongoing project in visual, numerical General Relativity at The Australian National University. Recently, the numerical differential geometric engine of *GRworkbench* has been rewritten using functional programming techniques. By allowing functions to be directly represented as program variables in C++ code, the functional framework enables the mathematical formalism of Differential Geometry to be more closely reflected in *GRworkbench*. The powerful technique of ‘automatic differentiation’ has replaced numerical differentiation of the metric components, resulting in more accurate derivatives and an order-of-magnitude performance increase for operations relying on differentiation.

1 Introduction

The goal of the ongoing *GRworkbench* project at The Australian National University is to create a visual software tool for numerical operations on analytically defined space-times in General Relativity. Such a tool would be used by researchers and educators alike to quickly gain insight into the physical properties of known exact solutions of the Einstein field equation.

A new version of *GRworkbench* was implemented in 1999 [1]. It featured a novel numerical differential geometric engine and a flexible visualisation system [2, 3], and was easy to extend with additional space-time definitions.

*andrew.moylan@anu.edu.au, susan.scott@anu.edu.au, antony.searle@anu.edu.au

Set	C++ type	Notes
\mathbb{Z}	int	max. $\pm(2^{31} - 1)$
\mathbb{R}	double	max. $\sim \pm 10^{308}$, precision 15
\mathbb{R}^n	<code>nvector<double></code>	(as for double)
$(A \rightarrow B)$	<code>function<B (A)></code>	see Section 2.1

Table 1: Correspondence between certain mathematical sets and C++ types in *GRworkbench*. The term ‘max.’ refers to the largest representable elements of the set; the term ‘precision’ refers to the number of significant figures to which elements of the set are represented.

In 2003, the numerical and differential geometric aspects of *GRworkbench* were rewritten using functional programming techniques, enabling the direct representation in the C++ code of *GRworkbench* of those concepts in numerical computation and Differential Geometry that are normally defined in terms of functions. The new functional framework enables potentially complex numerical experiments to be quickly and simply defined [4]. Numerical experiments performed in *GRworkbench* were employed in our analysis of a recent scientific claim, described elsewhere [5, 6].¹

Functional programming and its applications in *GRworkbench* are described in §2. Automatic differentiation, in which the derivatives of a function are exactly computed whenever the value of the function is computed, is described in §3. In §4 we outline planned future developments for the numerical engine of *GRworkbench*, in which interval arithmetic will be used to establish a guaranteed bound on numerical errors.

2 Functional programming

In the traditional programming languages of scientific computing, programs typically consist of *routines* that operate on data stored in program *variables*. Every variable in C++ has a *type*, and there is an approximate correspondence between C++ types and standard mathematical sets. Table 1 lists some important examples.

The first two sets in Table 1 are, of course, directly representable in some way in every language of scientific computing. The `nvector<T>` type uses the C++ template mechanism ([8], page 327) to provide a type representing n -tuples of any other type T . Thus, `nvector<double>` represents elements of \mathbb{R}^n and `nvector<nvector<double>>` represents elements of $\mathbb{M}_{m \times n}$, the set of $m \times n$ matrices with real-valued entries.

¹For the original scientific claim see [7].

The following is a routine in C++:

```
double mean(double a, double b)
{
    return (a + b) / 2;
}
```

The corresponding mathematical definition is

$$\begin{aligned} \text{mean}: \mathbb{R} \times \mathbb{R} &\rightarrow \mathbb{R}, \\ \text{mean}(a, b) &= \frac{a + b}{2}. \end{aligned} \tag{1}$$

The first line of the routine conveys the same information as the first line of (1): the routine `mean` takes two real numbers as arguments, and returns a real number. The first line of the routine without the routine name or argument names, i.e., **double (double, double)**, is the *signature* of the routine; it conveys the same information as $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ in (1). The rest of the routine definition, enclosed in braces, encodes the second line of (1).

We may define a *function* as anything which behaves like the routine `mean` above, in the sense that it accepts zero or more arguments, and returns a value. In the most common programming languages of scientific computation, C and Fortran, the *only* possible functions are routines, and so the terms ‘function’ and ‘routine’ are used interchangeably. The key feature of functional programming is that there can be functions other than the routines typed in by the programmer—functions created while the program is running. The mechanism to achieve this in C++ is introduced in §2.1.1; we first describe how functions can be stored in variables in C++.

2.1 Functions as data

The capability to store functions in variables is not unique to functional programming. Most languages used for scientific computation have some way to store a reference to a program routine. For example, in C and C++ the *address* `&f` of a routine **double f (double x)** can be stored in a *function pointer* variable of type **double (*)(double)**. Function pointers can be called just like routines.

GRworkbench uses the Boost Function library [9] to store references to functions. The Boost Function library provides the templatised type `function<T>` representing a function whose signature is `T`. The following code fragment shows how the routine `mean` can thus be stored in a variable:²

```
function<double (double, double)> f = mean;
// the following two lines are now equivalent
```

²Anything after the characters `//` in a line of code is ignored by the C++ compiler.

```

double x = f(1, 2);
double x = mean(1, 2);

```

Observe from the last two lines that the variable `f` can be used just like the routine `mean`; they are both functions.

In general, if we let $(A_1 \times \cdots \times A_n \rightarrow B)$ denote the set of functions from $A_1 \times \cdots \times A_n$ to B , then the corresponding C++ type in *GRworkbench* is `function<B (A1, . . . , An)>`, where the sets B, A_1, \dots, A_n correspond to the types `B, A1, . . . , An`. The fourth row of Table 1 summarises this relationship.

The most important application of the storage of functions in variables is that functions can then be arguments to other functions. Consider the following routine, which crudely approximates the derivative of a function f at a point x :

```

double slope(function<double (double)> f, double x)
{
    double h = 0.1;
    return (f(x + h) - f(x - h)) / (2 * h);
}

```

The corresponding mathematical definition is

$$\begin{aligned} \text{slope}: (\mathbb{R} \rightarrow \mathbb{R}) \times \mathbb{R} &\rightarrow \mathbb{R}, \\ \text{slope}(f, x) &= \frac{f(x+h) - f(x-h)}{2h}, \quad h = 0.1. \end{aligned} \quad (2)$$

Many other numerical algorithms naturally take a function as an argument. Two examples are

$$\begin{aligned} \text{minimise}: (\mathbb{R} \rightarrow \mathbb{R}) \times \mathbb{R} &\rightarrow \mathbb{R}, \\ \text{minimise}(f, x) &= (\text{a local minimum of } f \text{ near } x), \end{aligned} \quad (3)$$

and

$$\begin{aligned} \text{integrate}: (\mathbb{R} \rightarrow \mathbb{R}) \times \mathbb{R} \times \mathbb{R} &\rightarrow \mathbb{R}, \\ \text{integrate}(f, a, b) &= (\text{numerical estimate of } \int_a^b f(x) dx). \end{aligned} \quad (4)$$

2.1.1 Creating functions at run-time

Consider the following function, defined in terms of the slope function (2):

$$\begin{aligned} \text{derivative}: (\mathbb{R} \rightarrow \mathbb{R}) &\rightarrow (\mathbb{R} \rightarrow \mathbb{R}), \\ \text{derivative}(f) = g, \quad g: \mathbb{R} &\rightarrow \mathbb{R}, \quad g(x) = \text{slope}(f, x). \end{aligned} \quad (5)$$

Concept	Representation in <i>GRworkbench</i>
Coordinates	<code>nvector<double></code>
Metric components	<code>nvector<nvector<double>></code>
Inter-chart map	see Section 2.2.2
Point	<code>point</code>
Tangent vector	<code>tangent_vector</code>
Metric	<code>function<double (tangent_vector, tangent_vector)></code>
World-line	<code>function<point (double)></code>

Table 2: Representation of important differential geometric concepts in *GRworkbench*.

For each function f , it returns *the function which returns the slope of f at its argument*.

In most traditional languages of scientific computing it is not possible to encode a routine which returns `derivative(f)` for all functions $f: \mathbb{R} \rightarrow \mathbb{R}$. In C++ it *is* possible to encode this function, by means of a *functor class* ([8], pages 514–515). The result is a C++ routine `function<double (double)> derivative(function<double (double)> f)`. (For complete code, including that of the necessary functor class, see [6], pages 13–14, 92–94.)

If we were to replace the `slope` routine with a more sophisticated algorithm for numerical differentiation, then this `derivative` routine would be a good approximation to the mathematical operation of differentiation. For example, `derivative(sin)` would be a good approximation to the function `cos`.³ Functional programming permits numerical operations, like `derivative`, to be expressed in a way which closely resembles the mathematical operations that they approximate.

2.2 Functional Differential Geometry

Many fundamental notions in Differential Geometry and General Relativity, such as the action of the metric tensor, and particle world-lines, are functions. By elevating functions to the same level as traditional data types (\mathbb{Z} , \mathbb{R}), functional programming makes these notions directly representable as variables in the C++ code of *GRworkbench*. Table 2 summarises the correspondence between important concepts in Differential Geometry and their representations in *GRworkbench*.

³Many standard functions, including `sin` and `cos`, are built-in to C++.

2.2.1 Charts and the metric components

On the n -dimensional space-time manifold \mathcal{M} , a *chart* is a pair (\mathcal{U}, ϕ) representing a coordinate system on the set $\mathcal{U} \subset \mathcal{M}$, where the one-to-one function $\phi: \mathcal{U} \rightarrow \mathbb{R}^n$ maps points in \mathcal{U} to their coordinates in \mathbb{R}^n . Space-times are defined in *GRworkbench* by the specification of the components of the metric tensor on $\phi(\mathcal{U}) \subset \mathbb{R}^n$ for one or more charts (\mathcal{U}, ϕ) , and by the definition of the maps $\phi_\beta \circ \phi_\alpha^{-1}$ between the coordinate systems of overlapping pairs of those charts, where \circ denotes function composition.

The coordinates of a point on a chart, $x^i \in \mathbb{R}^n$, are represented by a variable of type `nvector<double>`. The components g_{ab} of the metric tensor at a point on a chart are represented as an $n \times n$ matrix, by a variable of type `nvector<nvector<double>>`. A function which defines the metric components as a function of the chart coordinates x^i might then be of the form

$$\begin{aligned} \text{chart}: \mathbb{R}^n &\rightarrow \mathbb{M}_{n \times n}, \\ \text{chart}(x^i) &= g_{ab}|_{x^i}, \end{aligned} \tag{6}$$

represented in *GRworkbench* by a function of signature `nvector<nvector<double>> (nvector<double>)`. In general, however, the chart coordinates are an open subset of \mathbb{R}^n , and so (6) will not be defined everywhere in \mathbb{R}^n . A mechanism is required to represent functions which are only defined on a subset of some other, standard set. (By ‘standard set’ we mean a set which is already represented by a type in C++, such as those listed in Tables 1 and 2.)

GRworkbench employs the *Boost Optional* library [10] to represent functions which are undefined for some values of their arguments. The Boost Optional library provides a templatised type `optional<T>`, which represents the set $S \cup \{\emptyset\}$, where S is the set corresponding to the template parameter type `T`, and \emptyset is a special value taken by functions at points where they are undefined.

Using the `optional` mechanism we rewrite (6) to support charts defined on subsets of \mathbb{R}^n :

$$\begin{aligned} \text{chart}: \mathbb{R}^n &\rightarrow \mathbb{M}_{n \times n} \cup \{\emptyset\}, \\ \text{chart}(x^i) &= \begin{cases} g_{ab}|_{x^i}, & \text{if the } x^i \text{ are valid chart coordinates;} \\ \emptyset, & \text{otherwise.} \end{cases} \end{aligned} \tag{7}$$

The corresponding C++ type is

$$\text{function<optional<nvector<nvector<double>>> (nvector<double>)>,} \tag{8}$$

for which *GRworkbench* declares the short synonym `chart`, using the C++ `typedef` mechanism:

```
typedef function<optional<nvector<nvector<double>>> (nvector<
double>)> chart;
```

The `optional` mechanism is most useful when the caller of a function cannot know beforehand whether the function will be defined at the arguments to be given to it. This is the case for callers of functions of the form (7). The `optional` mechanism thus enables the differential geometric and visualisation algorithms in *GRworkbench* to be coded in such a way that they can operate on any space-time definition, without prior knowledge of the particular coordinate systems (charts) in which they will be working.

2.2.2 Inter-chart maps

For two charts $(\mathcal{U}_\alpha, \phi_\alpha)$ and $(\mathcal{U}_\beta, \phi_\beta)$, the *inter-chart map* from $\phi_\alpha(\mathcal{U}_\alpha)$ to $\phi_\beta(\mathcal{U}_\beta)$ is

$$\begin{aligned}\phi_{\alpha\beta}: \phi_\alpha(\mathcal{U}_\alpha) &\rightarrow \phi_\beta(\mathcal{U}_\beta), \\ \phi_{\alpha\beta}(x^i) &= (\phi_\beta|_{\mathcal{U}_\alpha} \circ \phi_\alpha^{-1})(x^i),\end{aligned}\tag{9}$$

where $\phi_\beta|_{\mathcal{U}_\alpha}$ is the function ϕ_β restricted to the set \mathcal{U}_α .

The domain of $\phi_{\alpha\beta}$ is, in general, a subset of \mathbb{R}^n . Hence $\phi_{\alpha\beta}$ cannot be represented by a variable of type `function<nvector<double> (nvector<double>)>`; instead, the `optional` mechanism is again employed. Thus, an inter-chart map from a chart $(\mathcal{U}_\alpha, \phi_\alpha)$ to a chart $(\mathcal{U}_\beta, \phi_\beta)$ is represented by a function

$$\begin{aligned}\text{map}: \mathbb{R}^n &\rightarrow \mathbb{R}^n \cup \{\emptyset\}, \\ \text{map}(x^i) &= \begin{cases} (\phi_\beta|_{\mathcal{U}_\alpha} \circ \phi_\alpha^{-1})(x^i), & \text{if } (x^i) \in \phi_\alpha(\mathcal{U}_\alpha) \text{ and } \phi_\alpha^{-1}(x^i) \in \mathcal{U}_\beta; \\ \emptyset, & \text{otherwise.} \end{cases}\end{aligned}\tag{10}$$

The corresponding C++ type is

$$\text{function<optional<nvector<double>> (nvector<double>)>}.\tag{11}$$

The C++ `typedef` mechanism is used to define a synonym `map` for this type.

2.2.3 World-lines

The abstract notion of a point $p \in \mathcal{M}$, independent of any particular coordinate system, is represented in *GRworkbench* by a C++ type `point`. A `point` is constructed from two pieces of information: a `chart` which contains it, and its coordinates on that chart.

The **operator**[] routine of the **point** type takes one argument, a variable of type **chart**, and returns a variable of type **optional<nvector<double>>**, representing the coordinates of the point on the given chart. (The **optional** mechanism is used because a particular point may not have coordinates on the given chart.) Thus, if **p** is a variable of type **point**, and **c** is a variable of type **chart**, then the coordinates of **p** on **c** are given by **p[c]**. If there is no inter-chart map defined from the point's original chart to the chart **c** then \emptyset is returned.

A general curve in space-time, such as a world-line, which may not be defined for all values of its parameter, is a function $\lambda: \mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}$; such functions are represented by variables of type **function<optional<point>> (double)>**. Curves in *GRworkbench* must be represented in this form, because they are often defined in terms of numerical algorithms that may not converge everywhere to a solution, even if a solution exists; such a curve returns the value \emptyset for parameter values for which there was no convergence to a solution. The synonym **worldline** is defined for the type **function<optional<point>> (double)>**.

2.2.4 Tangent vectors

The abstract notion of a tangent vector $v \in T_p$, where T_p is the tangent space of a point $p \in \mathcal{M}$, is represented in *GRworkbench* by the C++ type **tangent_vector**. A **tangent_vector** is constructed from three pieces of information: the **point** to whose tangent space it belongs, a **chart** containing that point, and the contravariant components of the tangent vector on that chart.

As with the **point** type, the **operator**[] routine of the **tangent_vector** type takes one argument, a variable of type **chart**, and returns the components of the tangent vector on the given chart, in a variable of type **optional<nvector<double>>**. When the components of a tangent vector are requested on a chart other than that from which the tangent vector was constructed, *GRworkbench* uses the inter-chart map, if it exists, to compute the components. If v^i are the components of a tangent vector v at a point p on a chart with coordinates x^i , then the components on another chart, with coordinates $x^{i'}$, are

$$v^{i'} = \left. \frac{\partial x^{i'}}{\partial x^i} \right|_p v^i = A_i^{i'} v^i. \quad (12)$$

The entries of the matrix $A_i^{i'}$ are the derivatives of the inter-chart map $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^n$ with respect to the coordinates x^i of its argument, evaluated at p . In its most recent version, *GRworkbench* computes $A_i^{i'}$, and thereby the components $v^{i'}$, using the method of §3.

At a point p , the metric g_{ab} is naturally considered as the inner product

$$\begin{aligned} \text{metric}: T_p \times T_p &\rightarrow \mathbb{R}, \\ \text{metric}(u, v) &= g_{ab}u^av^b. \end{aligned} \tag{13}$$

This function is encoded in *GRworkbench* in the routine `metric`, whose signature is **double** (`tangent_vector`, `tangent_vector`). Also, the **operator*** routine of the `tangent_vector` type is defined to call `metric`, so that if `u` and `v` are variables of type `tangent_vector`, then the expression `u * v` is equivalent to the expression `metric(u, v)`. This notation is reminiscent of the two equivalent forms

$$g_{ab}u^av^b = u_bv^b \tag{14}$$

for the inner product of two tangent vectors.

2.3 Numerical operations and applications

We may think of the operation of numerically determining a geodesic by integrating the geodesic equation from initial conditions as a function

$$\begin{aligned} \text{geodesic}: T_p &\rightarrow (\mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}), \\ \text{geodesic}(v) &= \lambda, \quad \lambda: \mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}, \end{aligned} \tag{15}$$

where λ is the numerically determined geodesic with tangent vector v at $p = \lambda(0)$. Similarly, the numerical parallel transport of a tangent vector along a curve may be represented by a function

$$\begin{aligned} \text{parallel_transport}: (\mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}) \times T(\mathcal{M}) &\rightarrow (\mathbb{R} \rightarrow T(\mathcal{M}) \cup \{\emptyset\}), \\ \text{parallel_transport}(f, v) &= h, \quad h: \mathbb{R} \rightarrow T(\mathcal{M}) \cup \{\emptyset\}, \\ h(t) &\text{ is the parallel transport of } v \in T_{f(0)} \text{ to } f(t) \text{ along } f, \end{aligned} \tag{16}$$

or $h(t) = \emptyset$ if, for example, the numerical parallel transport algorithm did not converge to a solution.

In terms of these two functional definitions it is trivial to define the following interesting object:

$$\begin{aligned} \text{parallel_curve}: (\mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}) \times T(\mathcal{M}) &\rightarrow (\mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}), \\ \text{parallel_curve}(f, v) &= \gamma, \quad \gamma: \mathbb{R} \rightarrow \mathcal{M} \cup \{\emptyset\}, \\ \gamma(t) &= \text{geodesic}(\text{parallel_transport}(f, v)(t))(1), \end{aligned} \tag{17}$$

which represents the world-line of an object which is stationary at a proper distance $\sqrt{g_{ab}v^av^b}$ with respect to an observer whose world-line is f . By ‘gluing’ together functional objects such as these, it is easy to define potentially complex *numerical experiments* simulating interesting physical situations [5].

3 Automatic differentiation

Differentiation plays a role in most numerical operations in *GRworkbench*, because the derivatives of the components of the metric tensor feature in, for example, the geodesic equation and the parallel transport equation. Previously in *GRworkbench*, numerical differentiation of arbitrary functions $f: \mathbb{R} \rightarrow V$, where V is any vector space, was accomplished via numerical estimation of the limit

$$\lim_{h \rightarrow 0} \frac{f(x+h) - f(x-h)}{2h} = \lim_{h \rightarrow 0} d(h), \quad (18)$$

using the technique of *Richardson extrapolation* ([6], pages 18–21; [11], pages 186–189). In (18), $d(h)$ is the *centred difference* approximation to the derivative of f at x .

While this method is easily applicable to any function, it has some important drawbacks. Firstly, it requires many evaluations of the function f , at $x \pm h$ for numerous values of h . Secondly, an estimate of the limit (18) may not converge, or, what is worse, may converge to a value that is incorrect. The accurate convergence of the algorithm depends on the values of h at which $d(h)$ is evaluated being of approximately the same size as the smallest scale over which the function varies significantly around x . For a general implementation of the numerical differentiation method (18), this information may not be available. Thirdly, the accuracy of this method is rarely more precise than half as many significant figures as the precision of the floating point arithmetic employed. (For the **double** type this is $\sim 15/2 \sim 7$ significant figures.)

The technique of *automatic differentiation* ([12], Chapter 5) avoids all of these difficulties. The technique follows from three observations. The first is that all C++ routines encoded by the programmer must, at the lowest level, eventually be defined in terms of a finite set of built-in fundamental operations (such as addition, extraction of square roots, and trigonometric functions). The second is that all of these fundamental operations have exact derivatives at (almost) every point at which they are defined, and these exact derivatives may themselves be expressed in terms of the fundamental operations. The third is that from the exact derivatives of the fundamental operations we can obtain the exact derivative of any function defined in terms of the fundamental operations, through the use of the chain rule for differentiation:

$$\frac{d}{dx} f(u(x)) = \left. \frac{d}{du} f(u) \right|_{u(x)} \frac{d}{dx} u(x). \quad (19)$$

These facts are exploited in the method of *forward automatic differentiation*, which is implemented in *GRworkbench*. In this method, the

fundamental C++ operations which act on real numbers are extended to also act on ordered pairs (u, u') of real numbers, with u being the usual argument to the function, and u' representing the derivative of u with respect to some independent variable x . For example, we redefine

$$\begin{aligned}\sin: \mathbb{R} \times \mathbb{R} &\rightarrow \mathbb{R} \times \mathbb{R}, \\ \sin(u, u') &= (\sin u, u' \cos u),\end{aligned}\tag{20}$$

and

$$\begin{aligned}\text{multiplication: } (\mathbb{R} \times \mathbb{R}) \times (\mathbb{R} \times \mathbb{R}) &\rightarrow \mathbb{R} \times \mathbb{R}, \\ (u, u')(v, v') &= (uv, u'v + uv').\end{aligned}\tag{21}$$

In this way, the derivative of a function is automatically computed whenever the value of the function is computed. The derivative is accurate up to round-off error in the floating point arithmetic, and its accuracy is independent of the behaviour of the function around x . As well as being more accurate, the method is also faster than numerical differentiation, because the extra computation involved in computing the second part of each ordered pair (u, u') (roughly a factor of two) is generally less computationally expensive than the many evaluations of f that are necessary to numerically estimate the value of the limit (18).

The replacement of numerical differentiation by automatic differentiation has resulted in a large performance and accuracy increase for those numerical operations in *GRworkbench* that are dependent on the differentiation of the components of the metric tensor.

4 Future directions

Currently in *GRworkbench*, tangent vectors and metrics, respectively tensors of type $(1, 0)$ and type $(0, 2)$, are each implemented separately. We will replace them with an implementation of general tensors of arbitrary type, whose components on any chart are determined in terms of their components on one chart by the general transformation rules. This implementation will be based around a C++ type that can represent any function of vectors and covectors that is linear in each of its arguments. Thus, the representation of tensors in *GRworkbench* will be nearly identical to their usual mathematical definition.

4.1 Interval arithmetic

Many computations require us to make statements not about points, but rather about open sets of points. Ordinary differential equation solvers

require that functions be defined, bounded, and have some number of continuous derivatives not at a single point, but in a neighbourhood around an initial condition. Just as we can implement automatic differentiation by extending the definition of basic operations, we can also implement *interval arithmetic* ([12], Chapter 3), so that for

$$f: \mathbb{R} \rightarrow \mathbb{R} \quad (22)$$

we define a new function

$$f': \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R} \quad (23)$$

that satisfies

$$f'(a, b) = (c, d) \Rightarrow \forall x \in [a, b], f(x) \in [c, d]. \quad (24)$$

The function f' is a function from one interval of real numbers to another interval of real numbers. By combining the definition of a function in terms of fundamental operations with a bound on its parameter, we obtain a bound on the value of the function. For example, the function

$$\begin{aligned} \text{multiply}: \mathbb{R} \times \mathbb{R} &\rightarrow \mathbb{R}, \\ \text{multiply}(a, c) &= ac, \end{aligned} \quad (25)$$

is generalised to act on intervals of real numbers by

$$\begin{aligned} \text{multiply}' : (\mathbb{R} \times \mathbb{R}) \times (\mathbb{R} \times \mathbb{R}) &\rightarrow (\mathbb{R} \times \mathbb{R}), \\ \text{multiply}'((a, b), (c, d)) &= (\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)). \end{aligned} \quad (26)$$

Thus, the product of the interval (2.9, 3.1) with the interval (3.9, 4.1) is the interval (11.31, 12.71).

When combined with automatic differentiation and the **optional** mechanism, a whole range of statements about *any* function defined in terms of fundamental operations can be tested on any interval. We might be able to establish with certainty that a function is everywhere defined on a certain closed ‘box’ $[a, b] \times [c, d] \times \dots \subset \mathbb{R}^n$, that it is bounded, or even that it is C^k . These properties are important to guarantee that geodesic tracing and other algorithms can respond correctly to singularities and other pathological conditions.

A concrete example occurs for the metric components of the Schwarzschild space-time expressed in standard spherical polar coordinates, which are valid separately for $0 < r < 2M$ and for $2M < r < \infty$. Presently in *GRworkbench*, two **charts** are used, respectively covering $r \in (0, 2M)$ and $r \in (2M, \infty)$, so that a numerical operation, such as geodesic tracing, that

does not ‘know’ about the $r = 2M$ singularity, cannot accidentally ‘skip over’ it by discretely sampling points on either side of it.

Using interval arithmetic it would, in principle, be possible to employ one ‘pseudo-chart’ covering the whole range $0 < r < \infty$, because any numerical operation that evaluated the metric components on an interval containing $r = 2M$ would learn that the metric components are undefined somewhere in that interval, and are unbounded over it—the algorithm should then attempt to investigate a smaller subset of the original interval.

A more widely applicable further benefit of interval arithmetic is that it can provide powerful control over round-off errors. The exact result for any calculation typically lies between two representable floating-point numbers; round-off error occurs when one of these must be chosen to represent the result, a process analogous to writing a value to only so many decimal places. Instead of choosing one value, an interval arithmetic operation returns both values, which together bound the tightest representable interval enclosing the true result. Interval arithmetic trades the illusory precision of floating point computations for guaranteed accuracy.

Once it is implemented in *GRworkbench*, the ‘pseudo-algebraic’ method of interval arithmetic, which exploits exactly-known properties of the fundamental mathematical operations in C++, will enable a powerful new method of exploring properties of space-times and render the ability to make exact statements about these properties.

5 Conclusion

The new functional framework for the differential geometric engine of *GRworkbench* represents the underlying mathematical structure of Differential Geometry more closely than ever before. Functional programming makes it easier to define systems that model interesting physical situations in numerical experiments. The technique of automatic differentiation provides accurate and fast derivatives for functions, such as analytically-defined space-time metrics, that are defined in terms of certain fundamental mathematical operations. In the future, interval arithmetic may enable the numerical engine of *GRworkbench* to be written in such a way that there are no unknown numerical errors at all.

References

- [1] A. C. Searle. *GRworkbench*. Honours thesis, The Australian National University, 1999.

- [2] S. M. Scott, B. J. K. Evans, and A. C. Searle. *GRworkbench*: A computational system based on differential geometry. In V. G. Gurzadyan, R. T. Jantzen, and R. Ruffini, editors, *Proceedings of the Ninth Marcel Grossmann Meeting on General Relativity*, pages 458–467. World Scientific, 2002.
- [3] B. J. K. Evans, S. M. Scott, and A. C. Searle. Smart geodesic tracing in *GRworkbench*. *General Relativity and Gravitation*, 34:1675–1684, 2002.
- [4] A. Moylan, S. M. Scott, and A. C. Searle. Developments in *GRworkbench*. In *Proceedings of the Tenth Marcel Grossmann Meeting on General Relativity*, 2005. To appear.
- [5] A. J. Moylan, S. M. Scott, and A. C. Searle. Can the Milky Way be weighed using earth-based interferometry? 2005. In preparation.
- [6] A. Moylan. Numerical experimentation within *GRworkbench*. Honours thesis, The Australian National University, 2003.
- [7] M. Karim, A. Tartaglia, and A. H. Bokhari. Weighing the Milky Way. *Classical and Quantum Gravity*, 20:2815–2825, 2003.
- [8] B. Stroustrup. *The C++ Programming Language*. Addison-Wesley, third edition, 1997.
- [9] D. Gregor. *Boost Function library*, 2003. <http://www.boost.org/doc/html/function.html>.
- [10] F. Cacciola. *Boost Optional library*, 2004. <http://www.boost.org/libs/optional/doc/optional.html>.
- [11] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C; The Art of Scientific Computing*. Cambridge University Press, second edition, 1992.
- [12] R. Hammer, M. Hocks, U. Kulisch, and D. Ratz. Numerical toolbox for verified computing I. In *Springer Series in Computational Mathematics*, volume 21. Springer-Verlag, 1993.